Parametric Detail-on-Demand for Projection Spaces of High-Dimensional Data

Christian A. Steinparz[1,*], Yannick Metz[3,4], Andreas Hinterreiter[1], Manuela Waldner[2], Mennatallah El-Assady[3] and Marc Streit[1]

[1]Johannes Kepler University Linz, Linz, Austria
[2]TU Wien, Vienna, Austria
[3]ETH Zürich, Zürich, Switzerland
[4]University of Konstanz, Konstanz, Germany

Abstract

Inspecting the high-dimensional data underlying projection spaces requires an appropriate exploratory visual analytics approach. However, prior work often relies either on tedious manual selection of items and groups or on showing all details at once, which scales poorly and overwhelms users. We present an adaptive and adaptable visual analytics tool for sensemaking in complex high-dimensional spaces. Our tool provides a flexible set of parameters to dynamically determine what to show, where to show it, and at what level of detail. By combining user interaction with automatic interest computation and clustering, it produces continuously updated, map-like visualizations in which relevant structure is emphasized and contextual detail appears on demand across diverse datasets and domain-specific detail views. The interface is adaptive in how it updates annotation density, aggregation, and detail level, and adaptable in how it exposes user selections, queries, navigation, and parameter controls for steering this behavior.

Keywords

visualization, dimensionality reduction, dynamic labeling, multi-scale navigation, adaptable interfaces, adaptive interfaces

1. Motivation

Projection of high-dimensional data is widely used to create 2D representations of complex datasets. This process is also referred to as dimensionality reduction or embedding. The resulting 2D coordinates allow each data point to be visualized in a scatterplot. By aiming to preserve neighborhood relationships, such projections can make structures such as clusters or outliers visible [1, 2]. However, projections introduce two core challenges: (*i*) distortions can make spatial relationships unreliable, because distances in 2D may not faithfully reflect distances in the high-dimensional data; and (*ii*) scatterplots alone do not explain why clusters form or what they represent in the high-dimensional space [3, 4].

Annotations are a common solution, as they expose high-dimensional characteristics directly in the projection [3, 5]. However, current approaches tend to be either based on manual annotation [6, 7], requiring users to iteratively select and label clusters, which is slow and does not scale, or to show all details at once [8], which leaves information density uncontrolled and overwhelms users with visual clutter.

This leaves a gap for tools that support sensemaking of projection patterns by adaptively revealing relevant information as needed, offering adaptable behavior that can be configured to match user intent, and preserving navigational context so users stay oriented without cluttering the view.

Joint Proceedings of the ACM UMAP Workshops 2026, UMAP 2026, June 8–11, 2026, Gothenburg, Sweden
*Corresponding author.
christian.steinparz@jku.at (C. A. Steinparz); yannick.metz@uni-konstanz.de (Y. Metz); andreas.hinterreiter@jku.at (A. Hinterreiter); manuela.waldner@tuwien.ac.at (M. Waldner); elassady@ai.ethz.ch (M. El-Assady); marc.streit@jku.at (M. Streit)
0000-0002-6801-0438 (C. A. Steinparz); 0000-0001-5955-4487 (Y. Metz); 0000-0003-4101-5180 (A. Hinterreiter); 0000-0003-1387-5132 (M. Waldner); 0000-0001-8526-2613 (M. El-Assady); 0000-0001-9186-2092 (M. Streit)

2. Approach

We introduce a tool to augment projection spaces with adaptive, in situ annotations driven by Degree-of-Interest (DoI) [9, 10, 11]. Degree-of-Interest denotes a scalar value computed for each visual element to estimate how relevant or interesting it is expected to be for the user in the current interaction context. This value can then drive visual mappings, as higher values indicate that an item is currently more relevant to the user's analytic focus and should therefore be more likely to receive visual detail and emphasis. In our case, DoI determines the display state of points and clusters in the projection space: they may remain hidden, be shown without annotation, be annotated with text labels, or be annotated with detail-view insets. The tool offers a parametric interaction layer that allows users to control how this relevance is computed and propagated to other data through proximity in projection space and through data topology, how structure is aggregated, and how annotations are generated, directly shaping how details are revealed in the visualization.

2.1. Workflow

We assume a given 2D projection as input and use it as the spatial substrate for exploration. The tool is agnostic to the dimensionality reduction method and instead focuses on augmenting the resulting projection space with adaptive annotations. Projections with clear structure, such as well-separated clusters, provide the strongest basis for sensemaking.

During exploration, users *specify points of interest* through direct selections or feature search, defining an initial focus. For example, they may drag a lasso around an interesting cluster or enter feature expressions over data attributes and metadata, such as `material > 0.5 AND moves < 20` for projected chess positions. At this stage, the specified items form a hard selection: selected or queried items receive maximum DoI, while all other items receive zero DoI.

This hard focus is expanded via *DoI propagation*, turning the initial selection into a fuzzy relevance distribution. DoI spreads through proximity to neighboring points in the projection space, using spatial closeness as a proxy for data similarity, and through available data topology, such as trajectories in chess games, reinforcement learning episodes, or video sequences. As a result, DoI remains high in nearby dense clusters or along well-connected topology and decays for points that are farther away and less related.

The resulting DoI distribution drives *adaptive clustering*, organizing relevant elements into a hierarchical structure for multi-scale analysis.

Finally, *annotation and visual mapping* translate this structure into in situ labels and detail insets: visualizations of high-dimensional data features that reveal underlying characteristics directly within the projection and adapt continuously to pan-and-zoom interactions, similar to multi-scale map navigation. In this interaction loop, users act by specifying the initial focus, navigating the projection through pan and zoom, or adjusting parameters, while the tool adapts the DoI distribution, clustering, annotations, and layout in response.

2.2. Configuration

The described behavior is adapted through the tool's parametric interaction layer, allowing users to make the process more specific to their analytic intent. Users can steer how DoI propagates, for example by disabling spatial proximity propagation when they only want relevance to follow selected topologies. They can also regulate information density by controlling how much DoI is required before an item is annotated, how many annotations are shown, and whether the layout should fill available whitespace or remain sparse enough to preserve overview. Additional parameters control where annotations are placed, balancing proximity to their associated clusters against avoiding overlap with the scatterplot. Finally, basic visual mappings, such as the feature used for color encoding or text labels, can be adjusted to match the current analysis goal.

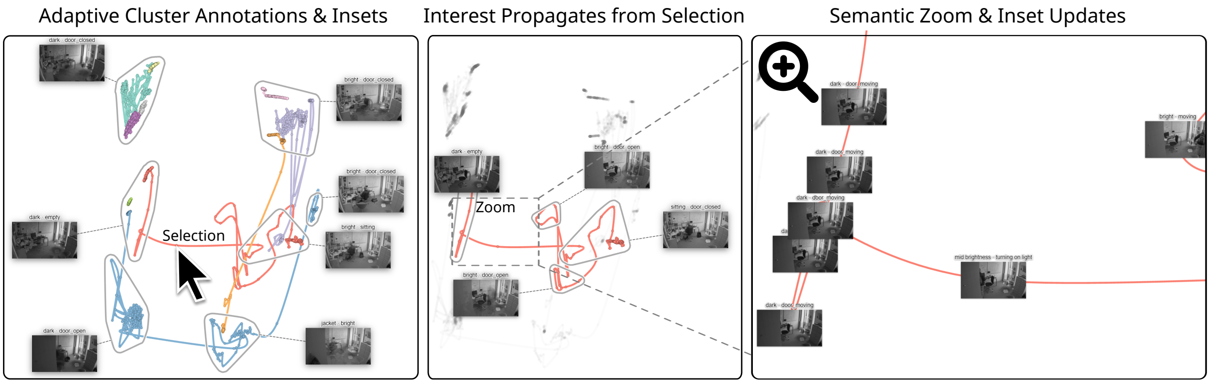

Figure 1: Adaptive annotation of a CCTV trajectory embedding: automatic cluster insets reveal structure (left), selection and resulting DoI propagation reveal annotations on a selected trajectory (center), and semantic zoom exposes details (right).

3. Implementation

The tool is implemented entirely in the frontend using React and TypeScript. Rendering is split between a custom WebGL2 pipeline for points and trajectories and JSX-based components for annotations and insets.

To support real-time interaction, we applied various optimizations: DoI propagation is approximated via neighborhood graphs with incremental updates. For clustering, we precompute an HDBSCAN hierarchy and refine interaction-driven runs asynchronously in a Web Worker. Annotation placement builds on a simulated annealing approach [12] with caching, hysteresis, and staged animations.

4. Use Case

Figure 1 illustrates our approach applied to a dataset of CCTV recordings of an office environment [13], thereby demonstrating one of many possible dataset-specific detail-view configurations. Each day of recordings forms a trajectory through the embedding space. Without annotation, users would see only an unlabeled scatterplot of points and edges.

Our tool automatically identifies the most salient clusters given the current zoom level and renders aggregate insets: visualizations summarizing all samples within a cluster rather than individual frames (Fig. 1, left). The embedding reveals a clear structure: the left column corresponds to dark rooms and the right to bright rooms, with closed-door states toward the top and open-door states toward the bottom, and two salient transitions between columns.

When a user selects a sample of interest, DoI propagates along the trajectory topology (Fig. 1, center): annotation layout dynamically refocuses to show insets only for clusters along the relevant trajectory. Zooming into the light-to-dark transition (Fig. 1, right) triggers semantic zoom: insets progressively update to finer-grained clusters and eventually to individual samples. This reveals how the door progressively opens along the trajectory and how the embedding makes a sharp horizontal jump at the moment a light is switched on.

In this use case, behavioral traces in the office recordings can reveal patterns such as room usage, routines, transitions, or anomalies. Game-related embeddings such as chess openings could similarly support player modeling by revealing strategic preferences or recurring decision patterns. Related projection-based workflows have also been used to analyze provenance data and interaction histories [14].

Acknowledgements

The work was partly funded by the Austrian Science Fund under grant number FWF DFH 23–N.

Declaration of Generative AI

The authors used OpenAI's ChatGPT (GPT-5) via https://chatgpt.com/ and Grammarly via https://www.grammarly.com/browser/chrome as writing assistants. The tools were used in an iterative human-in-the-loop manner to suggest rephrasings and improve grammar and structure, similar to a rubber-duck debugging process for prose. All AI-generated text was critically reviewed, edited, and rewritten by the authors. No generative AI tools were used to create or modify data or figures.

References

[1] L. v. d. Maaten, G. Hinton, Visualizing data using t-sne, Journal of machine learning research 9 (2008) 2579–2605.

[2] L. McInnes, J. Healy, J. Melville, Umap: uniform manifold approximation and projection for dimension reduction. arxiv, arXiv preprint arXiv:1802.03426 10 (2018).

[3] L. G. Nonato, M. Aupetit, Multidimensional projection for visual analytics: Linking techniques with distortions, tasks, and layout enrichment 25 (2019) 2650–2673. doi:10.1109/TVCG.2018.2846735.

[4] J. Stahnke, M. Dörk, B. Müller, A. Thom, Probing projections: Interaction techniques for interpreting arrangements and errors of dimensionality reductions 22 (2016) 629–638. doi:10.1109/TVCG.2015.2467717.

[5] M. D. Rahman, B. Doppalapudi, G. J. Quadri, P. Rosen, A survey on annotations in information visualization: Empirical insights, applications, and challenges 2410.05579 [cs.HC] (2024). doi:10.48550/arXiv.2410.05579.

[6] K. Eckelt, A. Hinterreiter, P. Adelberger, C. Walchshofer, V. Dhanoa, C. Humer, M. Heckmann, C. A. Steinparz, M. Streit, Visual exploration of relationships and structure in low-dimensional embeddings 29 (2023) 3312–3326. doi:10.1109/TVCG.2022.3156760.

[7] A. Hinterreiter, C. A. Steinparz, M. Heckmann, H. Stitz, M. Streit, Projection path explorer: Exploring visual patterns in projected decision-making paths 11 (2021) Article 22. doi:10.1145/3387165.

[8] S. Raval, C. Wang, F. Viégas, M. Wattenberg, Explain-and-test: An interactive machine learning framework for exploring text embeddings, 2023, pp. 216–220. doi:10.1109/VIS54172.2023.00052.

[9] G. W. Furnas, The FISHEYE view: A new look at structured files, Technical memorandum, Bell Laboratories, New York, 1981.

[10] F. van Ham, A. Perer, "search, show context, expand on demand": Supporting large graph exploration with degree-of-interest 15 (2009) 953–960. doi:10.1109/TVCG.2009.108.

[11] J. Abello, S. Hadlak, H. Schumann, H.-J. Schulz, A modular degree-of-interest specification for the visual analysis of large dynamic networks 20 (2014) 337–350. doi:10.1109/TVCG.2013.109.

[12] F. Lekschas, M. Behrisch, B. Bach, P. Kerpedjiev, N. Gehlenborg, H. Pfister, Pattern-driven navigation in 2d multiscale visualizations with scalable insets 26 (2020) 611–621. doi:10.1109/TVCG.2019.2934555.

[13] T. Qasim, R. B. Fisher, N. Bhatti, Ground-truthing large human behavior monitoring datasets, in: 2020 25th International Conference on Pattern Recognition (ICPR), IEEE, 2021, pp. 2763–2770.

[14] C. Walchshofer, A. Hinterreiter, K. Xu, H. Stitz, M. Streit, Provectories: Embedding-based analysis of interaction provenance data, IEEE Transactions on Visualization and Computer Graphics 29 (2023) 4816–4831. doi:10.1109/TVCG.2021.3135697.

A. Online Resources

An interactive demo of the tool is available at https://jku-vds-lab.at/paraadapt26-steinparz/.